

Use-Case in Delta Learning

Generative Adversarial Networks (GANs) are widely used for **domain transfer**, e.g. CycleGAN, and **data-augmentation** with controllable image synthesis, e.g. Generative Radiance Fields (GRAF). This makes them a fundamental tool for closing the domain delta in existing datasets. However, GANs are **challenging to train**: They need careful regularization, vast amounts of compute, and expensive hyper-parameter sweeps. Our approach achieves significant headway on these issues **improving image quality, sample efficiency, and convergence speed**.

Technical Problem

In GAN training, the discriminator aims to distinguish real from fake samples. On closer inspection, the discriminator's task is two-fold: First, it projects the real and fake samples into a meaningful space, i.e., it learns a representation of the input space. Second, it discriminates based on this representation yielding the training signal for the generator. In our work, we explore the utility of **pre-trained representations** to facilitate the discriminator's task and improve and stabilize GAN training.

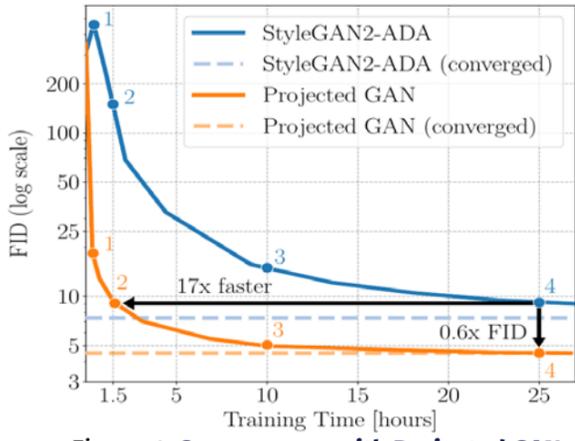


Figure 1: Convergence with Projected GANs. Our approach significantly speeds up convergence while yielding lower FIDs.

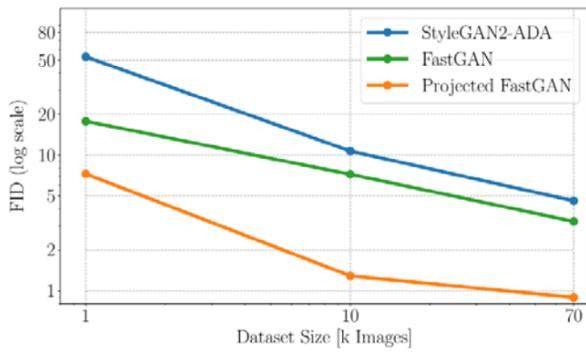


Figure 4: Small datasets. Projected GANs significantly improve FID scores, even for subsets of CLEVR with 1k and 10k samples.

Technical Solution

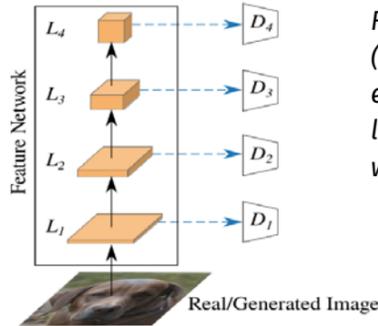


Figure 2: CCM (dashed blue arrows) employs 1x1 convolutions with random weights.

Cross-Channel Mixing (CCM). Empirically, we found two properties of the projection to be desirable: (i) it should be information preserving (ii) it should not be trivially invertible. The easiest way to achieve this is by mixing across channels with random 1x1 convolutions.

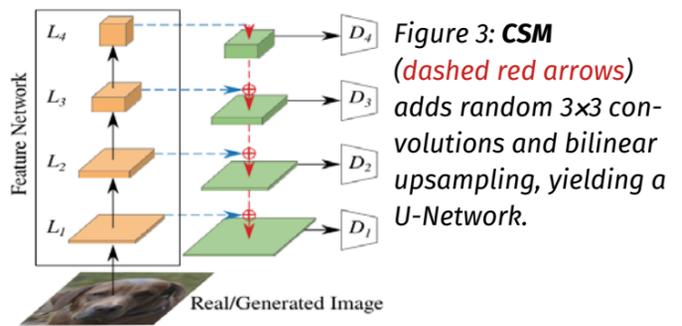


Figure 3: CSM (dashed red arrows) adds random 3x3 convolutions and bilinear upsampling, yielding a U-Net.

Cross-Scale Mixing (CSM).

To encourage feature mixing across scales, CSM extends CCM with random 3x3 convolutions and bilinear upsampling, yielding a U-Net architecture.

Evaluation

In our ablations on various pre-trained feature networks, we find that EfficientNets outperform both ResNets and Transformers. We thus use **EfficientNet-Lite1** as our feature network. Our method achieves **SOTA performance on both large and small scale datasets**.

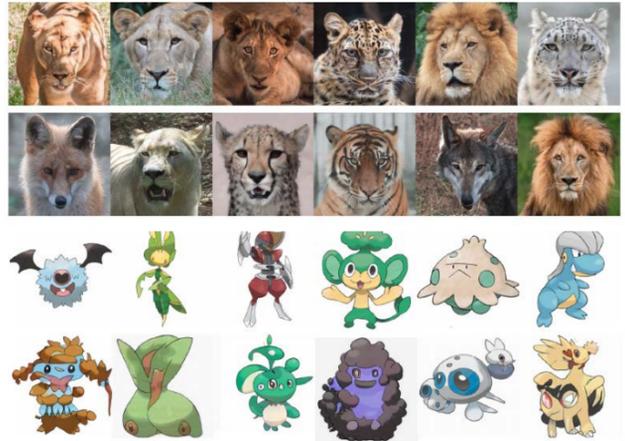


Figure 5: Real samples (top rows) vs. samples by Projected GAN (bottom rows) for small datasets with 4812 and 833 images, respectively.

Partners



External partners

